ORiON, Vol. 2, No. 2, pp. 34-51 *ISSN 0259-191X*

THE USE OF VISUAL INTERACTIVE SIMULATION TECHNIQUES FOR PRODUCTION SCHEDULING

W H SWANN
Imperial Chemical Industries PLC
Corporate Management Services Department
Fulshaw Hall
Wilmslow
Cheshire SK9 1QB England

ABSTRACT

During the last decade visual interactive simulation has become established as a useful new tool for solving real life problems. It offers the Operational Research professional the opportunity to impact beneficially on important new decision making areas of business and industry. As an example, this paper discusses its application to the scheduling of production on batch chemical plants, which to date has remained largely a manual activity. Two different approaches are introduced, and it is concluded that while discrete event simulation is most useful as an aid to learning at a time of change, bar chart simulation is preferred for the day to day scheduling. The technique has been implemented on a number of plants and has led to significant improvements in their performance. Some areas for further development are identified.

1. INTRODUCTION

Simulation is well established as an effective aid to decision-making in business and industry across a range of different applications. During the 1970s advances in computing technology opened the way for new, imaginative implementations of such aids which seemed to offer the opportunity to increase their impact. In particular the technique of visual interactive simulation began to emerge and has been readily embraced by practitioners to considerable effect (Bell, [1]).

The concept underlying the technique is to produce a system which a decision-maker can use directly himself, instead of via third parties as is the case with conventional models, to explore alternative scenarios and identify and assess possible courses of action. Easy to use interactions give the user full control over the model, and colour graphic displays linked to the model provide him with ready feedback on the consequences of potential decisions.

The approach helps improve the effectiveness of simulation in existing application fields, and provides the opportunity to extend its application to other problem areas hitherto relatively untouched by management science. This paper introduces visual interactive simulation in terms of its use in the scheduling of production in multistage, multiproduct, batch chemical plants as an example of the latter.

2. PRODUCTION SCHEDULING IN BATCH PLANTS

Batch production is used widely throughout the chemical industry, particularly in the manufacture of fine chemicals such as dyestuffs, pharmaceuticals etc. Products are made as discrete batches which progress through a pre-determined sequence of processing stages, each of which might be carried out in a different plant item of equipment. Over recent years market pressures have led to an increasingly customised product range offering a large number of variants, so the plants must have a multipurpose capability, while competitive pressures place a heavy premium on a rapid response to provide good customer service, giving extra emphasis to the need for flexibility. In addition financial pressures impose the requirement to maximise the utilisation of resources (plant items, manpower, materials, services) and at the same time minimise inventory costs. Thus the efficient management of batch plants is increasingly important to a wide sector of the chemical industry.

However the scheduling of production on such plants to exploit effectively their flexibility is a skilled activity. A typical plant may consist of say 5 process stages (eg heating, reacting, mixing, packing, etc) with 4 or 5 alternative plant items at each stage.

Each process operation will draw on some combination of resources which of course have a finite availability. The product range may cover 100 variants (with an annual turnover of perhaps £40 million) manufactured via different combinations of items and stages with several different possible routes through the plant for each of them. There may be 20 or more batches in progress at any one time covering up to 10 different products. Intermediate storage might be provided between stages, and batches may split and/or merge as they proceed through the plant. Working patterns may mean that the plant does not operate at nights or weekends, except under special circumstances, which limits the times at which batches can enter certain stages. Thus the combinatorial complexity can be considerable with an intricate network of possible routes through the plant, each node of which is a potential decision point, and batches competing with one another for access to the limited available resources.

Decisions concern the order in which products are manufactured, assignment of plant items to batches, allocation of the other resources, and timing of production. They are made in the context of several objectives which cover:-

- business requirements primarily concerned with achieving a desired level of customer service, which is likely to vary from product to product, from customer to customer, and from time to time;
- production requirements which demand safe operation of plant under stable patterns, minimising losses due to down-time at changeovers between products, absorbing quickly but smoothly the potentially disruptive consequences of disturbances to plans, while maximising the utilisation of plant and other resources, and maintaining good labour relations;
- working capital requirements which generally restrict stock levels to lie within limits imposed by senior management in the light of what they believe the business can afford.

Clearly these objectives, rather than being mutually independent, will often be conflicting so that scheduling becomes a balancing act seeking a suitable compromise between them in the context of the business policy. However this latter is unlikely to be a formally written set of rules. It will tend to be an amalgamation of the undocumented and unstructed set of views, attitudes and goals of the management group, which will fluctuate according to a range of influential factors such as the economic environment. Thus the difficulties resulting from the computational and logical aspects of scheduling are exacerbated by political issues which result from attempting to reconcile and integrate the desires and needs of different parts of the organisation, eg marketing v production v maintenance.

Hence generating an acceptable schedule can be an extremely difficult task. Attempting to assess the consequences of disturbances caused eg by priority changes or sudden resource non-availability, and to modify the schedule to cater for these is often many times worse!

The richness of the variety inherent in this situation means that a scheduler is encountering novel problems daily, and their complexity and the non-routine manner in which they are handled makes it very difficult to define a set of rules which accurately captures the decision process. An analogy with chess seems relevant in that it is very difficult to specify what makes a good schedule (move), but a skilled planner (grand-master) can recognise one when he sees it.

The decision making process is heavily dependent upon the human creative ability to discover attractive candidates for evaluation, to recognise patterns formed by the decision variables and their implications, and to exercise judgement in assessing the relative merits of alternatives. Because there is no definable right way to make a decision, two schedulers faced with the same problem will often adopt quite different approaches and will arrive at conclusions based on value judgements from their own personal frame of reference.

Although a significant effort has been devoted by operational researchers into the scheduling of job-shops, it has had little impact on the associated problem of batch chemical plants (Rippin, [9]). This is because the techniques developed offer little or no scope for human intervention as they progress towards a solution. They tend to optimise independently some conceptual abstraction of the given problem. Since the solution then appears as if from a black-box, the scheduler often has no feel for how he might tune it to allow for the judgemental and often fuzzy considerations not catered for by the technique employed but essential for acceptance and practical implementation. Therefore scheduling of multipurpose batch plants remains by and large a manual activity.

3. DECISION SUPPORT SYSTEMS

If computers are to be used to help more effectively in this field then a system is needed which forms an environment whereby computer and scheduler can work directly together to arrive at a good practical solution. The computer should carry out the structured elements of the task of solving problems by exploiting its computational, memory, and representational powers, and provide the framework within which the user can tackle the unstructured elements using his judgemental, intuitive, creative and pattern recognition powers.

Such systems are referred to as decision support systems and provide for:

- the decision maker to use his own skills and attributes to identify tentative solutions;
- the computer to handle the data management and calculation necessary to evaluate each of these;
- the decision maker to weigh up the evaluations in terms of the management objectives and policies to arrive at a decision.

This is likely to be an iterative process, so such systems are best implemented as interactive programs designed for direct hands-on use by the decision maker himself and under his control at all times.

Decision support systems are not concerned with imposing solutions by automating the decision process through predefined objectives. Rather the aim is to allow the decision maker to employ whatever decision process he chooses and provide him with very rapid feedback of the consequences of particular actions. This should enable him to increase his understanding of the structure underlying the problem situation so that he can improve his effectiveness in dealing with it.

During the 1970s there was a significant growth in the availability of relatively low cost, readily accessible processing power in the form of mini- and micro-computers. This combined with the increased functionality being provided in visual display unit terminals stimulated interest in extending existing approaches to modelling. In particular R D Hurrion took advantage of these developments to implement decision support systems in the form of visual interactive simulation models. It is from this root that much of the current activity in this area in the UK has grown.

Hurrion's work was initially based on the application of discrete event simulation to job-shop scheduling (Hurrion, [8]), but it has subsequently been extended with effect to the scheduling of multipurpose batch chemical plants.

4. DISCRETE EVENT SIMULATION

The basis of visual interactive modelling is to provide, on the vdu screen, a suitable representation of the problem under investigation linked to the underlying model being used. When the latter is invoked it directs information on its current state to the screen in a dynamic way so that the representation picture can be updated to keep the user informed about progress. A range of interactions is also provided to enable the user to interrupt a run of the model on demand so that he may redirect the run, request additional information from the system, or provide the model with new data before continuing.

In the case of applying discrete event simulation to the batch plant scheduling program, the most useful representation is usually a diagram which mimics the target plant, in either spatial or logical terms, whichever the scheduler finds conceptually more convenient. As a simulation proceeds through time, at each event the screen picture is updated to display the current plant conditions using an appropriate combination of alphanumerics, colours, icons, animation etc. Part of the display is dedicated to showing, often in graphical form, one or more of the relevant performance measures such as output volume, stocks, costs, resource utilisations, timeliness of products etc.

Thus a scheduler using the system is effectively provided with a silent film of the operation of his plant, together with an indication of how well it is performing against the chosen criteria. At any stage of a run he can interrupt the simulation and freeze time while he requests displays of different performance parameters, inputs new conditions, constraints etc, or imposes a plant operating decision on the system, before continuing the run. In this way he gains immediate feedback on the effect of any input decisions, and so can explore various strategies and their effect on the chosen performance measures to determine the most suitable one to adopt. The system thus provides a useful means of "what if" experimentation for the user.

The benefits of this approach fall into three main categories.

The model building process itself benefits. The display should be in the language and terminology of the scheduler, not that of simulation or some computer jargon. Thus he can be involved directly in the model development from a very early stage, which greatly improves the processes of verification and validation. Furthermore, since he is in effect part of the total model himself at run time, the computer model does not need to be complete in every detail such that it contains logic to cover every possibility no matter how remote. Rare situations can be handled much more effectively by user interaction. The results of these considerations is a significant improvement in the quality and speed of the modelling process, and a major increase in confidence in the final result on the part of the user.

The use made of the model benefits. The increase in user confidence usually brings with it a greater willingness and ability to exploit the model fully and to accept the schedules generated by it. In addition the modus operandi of direct use provides the scheduler not only with knowledge of the final results of a set of decisions, but insight into how they came about through a greater understanding of the relevant dynamics. For example it becomes much easier to identify where bottlenecks occur on the plant and the circumstances which lead to them. This in turn helps not only the decision-making of immediate concern, but that of the longer term too as an improved rationale develops as a framework for the future.

Finally commitment to the resulting schedule benefits. The improved rationale, together with the visual nature of the system, can be used in communicating the schedule to plant operating personnel and the result is generally a significant improvement in their understanding and thence their commitment to it. It has already been indicated that one of the difficult tasks of the scheduler is to cater for the different viewpoints of the various interested functions such as production, maintenance and marketing. A model such as that described makes it easier to explore those different viewpoints and to demonstrate their consequences. This usually leads to a greater mutual understanding between the functions and improved integration and/or reconciliation of those viewpoints. Decision-making becomes more rational and less emotive.

The major disadvantage of the approach is that at any point during a run the display only gives a snapshot view of the situation at the current simulation time, perhaps supported by a history if this has been recorded. There is no forward view. Thus decisions about the ordering of production and the allocation of resources have to be made when an opportunity to start production arises without any guidance on future constraints. For example there is no guarantee that product due dates, required to provide the desired level of customer service, can be met. There is no direct information on the knock-on effect a decision with respect to one batch will have on other batches being progressed through the plant until simulation time proceeds and the effects come into play. It may then be desired to revise an earlier decision, but this is impractical without restarting the run.

The exploration of alternatives is also rather clumsy. It is simple to reset the model to a particular decision point to restart a run if the data defining its state was stored when that point was reached initially. A revised decision can then be applied, but all the subsequent decisions have to be re-input as the simulation progresses unless some regenerative recording facility has been provided. It is also difficult to compare readily two different sets of decisions or schedules, although comparison of their end points is easy.

Furthermore scheduling is an on-going process which typically must cope with frequent disturbances resulting from breakdowns, cancellation of or new priority orders, delays in resource availability etc. The effects of these on the existing schedule must be assessed and a revised scheme devised to deal with the altered circumstances. This can be done using simulation, but the above disadvantages again apply.

5. BAR CHART SIMULATION

These problems can be overcome by replacing the discrete event simulation algorithm within the visual interactive model with one based on the use of bar charts. To exploit this the mimic diagram on the screen is replaced by a display of the schedule in the form of a Gantt chart in which the horizontal axis represents time and each division on the vertical axis represents a critical resource. The use of a particular resource in the schedule is represented by the presence of a bar or coloured strip on the chart in the relevant resource division and over the appropriate time period. This of course emulates the planning boards seen in many scheduler's offices. It has the considerable merit of displaying the complete schedule over a given time horizon so that its global characteristics (product availability, stocks etc) can be readily inferred.

The model supports the scheduling process by displaying the current plan over the relevant time horizon, again in combination with selected performance measures. The progress of any single product through the plant is stored as a network of occupations on the resources by the relevant batches and is depicted by the corresponding set of bars.

Interactions are provided to update a schedule by introducing a new network onto the existing schedule through allocating bars onto the Gantt chart at the chosen positions. For any product it is often possible to reduce the updating to allocation of a single bar against the most critical resource with the remaining bars in the network, both downstream and upstream, being generated automatically according to logic agreed by the scheduler and written into the model. Special cases are catered for by use of interactions designed to override the default logic.

Further interactions are provided to input disturbances which generally involve a process stage in the plant taking a different time to that planned (change length of bar), or being carried out at a different time (time shift bar), or using a different resource (reallocate bar), or some combination of these. Depending on the degree to which batches share common resources and/or compete for scarce ones, displacement of a single bar can have substantial knock-on effects on many of the other batches progressing through the plant. On many plants these can be extremely difficult to identify from the conventional planning board. However if the model is set up to maintain integrity of the batch networks at all times, then it can automatically move bars representing the affected batches so that the effects can be seen immediately from the Gantt chart.

The whole schedule can be viewed at any time, changes to the schedule are easily applied with ready feedback, and of course comparison of alternatives is easily achieved eg by displaying them on different lines on the Gantt chart. Hence this approach readily overcomes the difficulties identified with discrete event simulation.

Furthermore it clearly retains most of the benefits of the latter as identified above and in some ways enhances them. For example the Gantt chart representation is quickly recognised and understood by the plant operators who have to implement the schedule. This reinforces the communicability feature, and the current chart can be left on permanent display for ready reference.

Since the schedule is stored electronically, modern communication facilities enable it to be transmitted virtually anywhere using the appropriate equipment. This facility is most useful for businesses operating several plants geographically dispersed. The centre can readily view the various schedules being adopted and thus perform its co-ordinating role more effectively, and it can also use the system for its more strategic planning activity. Again this leads to greater mutual understanding and co-operation and thence improved performance.

The bar chart model is much easier to update than the discrete event based one, which makes it more effective as an aid to management control. This feature also helps scheduling which of course is an ongoing incremental activity and generally involves rolling the existing schedule forward through time while adding batches rather than creating a totally new one from scratch.

The main disadvantage is that the approach is much less effective for developing the appreciation of the underlying dynamics which tends to emerge so clearly from the discrete event method. For most of the time in a stable situation this is not a particularly damaging drawback. However when a significant change occurs, eg in plant management, in plant configuration, or in product range, dynamic considerations become important and the discrete event approach is then very useful as an experimental vehicle for accelerating progress along the learning curve.

In addition since the allocation of bars is generally done against the primary resources, it may be difficult to handle secondary resources and maintain their integrity.

Finally there may be difficulties if large amounts of intermediate storage are used in the plant. Since this provides a buffering or decoupling function, then when attempting to assess the knock-on consequences of a disturbance it can be very difficult to know how to deal with batches downstream of the storage. In most cases this is best left to the scheduler to cope with using suitable interactions.

Overall the general conclusion tends to be that the bar chart approach is preferred for the day to day scheduling.

6. BENEFITS

The use of visual interactive modelling in scheduling has been developed within ICI since the late 1970s and has been implemented successfully at a number of locations, eg Bowen et al [2], da Silva et al, [3]. Evaluations carried out by user management demonstrate that significant benefits have been gained.

Step changes in plant throughput of the order of 20% or more are reported.

Reductions of up to 30% in raw materials, work in progress, packages, and finished product inventory have been achieved.

Customer service levels have been at least maintained and in many cases improved.

These result from:-

- better scheduling, through generation and comparison of alternatives and the assessment of schedules against agreed criteria;
- better management control, through effective handling of disturbances;
- more stable schedules leading to more predictable deliveries, because of the greater integration and mutual understanding of the different business functions;
- better teamwork, because of the mutual understanding leads to more co-operative working and less stress and frustration;
- great commitment to schedules by staff at all levels who have a clearer view of objectives and targets and an improved understanding of the implications of deviation from plan.

7. IMPLEMENTATION ISSUES

The development of visual interactive systems is a skilled activity. It demands not only a high level of modelling ability and technical knowledge, but also interpersonal qualities and sensitivity in dealing with people.

The system should contain the degree of sophistication appropriate to the problem situation and to the user's knowledge, experience and abilities. Thus it is critical that users are involved in the design and development states, but it is important to recognise their shortcomings as system designers. Managers are not used to inventing computer systems, but once they are experienced users they are generally very able and willing to identify desirable new features. A prototyping approach is most effective, putting something meaningful into the hands of the user as quickly as possible and allowing the system to evolve from that point on. Design and implementation are inseparable and evolutionary, but the evolution must be a conscious strategy and not open-ended and unplanned.

This form of approach emphasises the need to build a modular system with the model, display, dialogue and data management sections developed as separate sub-systems linked by common data structures.

The model itself should be adaptive, simple, but not simplistic, complete on the important issues, easy to control, and conceptually accessible to the user. It should cope automatically with the decisions required to produce a schedule by providing default logic and rules, but must provide access to these for the dialogue sub-system so that the user can override them and impose his own decisions. However this must be done in a robust way so that the user cannot leave the model in an unrealistic state.

To the user the interface is the system and it must therefore score well on behavioural characteristics. The displays should be tailored to the user's requirements, with colour and movement used to highlight points of interest or concern. In particular the layout should match the scheduler's mental image of the plant as this will enhance his ability to recognise patterns as they emerge. Care needs to be taken not to clutter the screen with extraneous factors - only information relevant to the issue being considered should be displayed.

This can be achieved by forming a large number of suitable logical displays and allowing the user to map any combination of these onto the physical screen according to his requirements.

The interactions should be simple to use and remember, and should be robust and forgiving in that they should tolerate input errors and not easily crash the system. They should lead the user into a dialogue which guides him through the decision process, prompting and steering him from time to time, but ultimately leaving control very clearly in his hands.

As with any system a visual model is only as good as the data provided to it. Experience shows that the data needed for the scheduling task is not always readily available! If it is, and even if it is already computerised, it is unlikely to be in a form suitable for direct use in the model. It is very easy to underestimate the considerable effort that can be required to integrate the model into the appropriate data systems such that the relevant data is collected, maintained, and provided in suitable form.

Introduction of a scheduling system as outlined above is certain to have a significant effect on the working methods of the staff involved. It will help to formalise the decision process and begin to increase the degree of structure in it by making more explicit the application of judgement, whilst increasing the profitable communication around the business. However a system which cuts across organizational groupings, intrudes on territorial rights, or is inconsistent with the organization's structure and lines of communication, stands little chance of success. Therefore details of the design must match its organizational context in terms of patterns of access to and control over data, allocation of authority, responsibility for performance evaluation and action taking.

The final issue concerns the use of the system, with particular reference to the exploration of alternatives. Often schedulers are under such pressure that they merely search for a feasible solution and having found one immediately proceed to implement it. A system used in this way might well improve their efficiency, but is unlikely to impact on their effectiveness and as such may well not pay for itself. At best the schedulers might adopt an incrementalist attitude by limiting their analysis to alternatives which differ only marginally from the existing schedule. This may be considered a remedial approach, moving away from ills rather than towards predetermined and agreed objectives. As such it reduces the potential for gaining significant benefit under normal conditions, and is likely to prove ineffective in periods of substantial change when a major rethink is needed. Thus it is very important to work closely with the users during their early experience with the system to encourage them into good habits and ways of working so that they exploit the system to the full.

8. FUTURE DEVELOPMENTS

Although major advances have been made in recent years in the application of visual interactive simulation to scheduling, there remain opportunities for development in several areas which could yield further desirable increases in the effectiveness of such systems.

One such area concerns the technical facilities offered by the system. It is clear that the system should not impose an alien planning method on the scheduler, but should mesh with his existing activities. This of course is one of the benefits of the bar chart technique in that it aligns closely with many approaches currently in use. To be effective the system must provide a wide range of very basic interactions which the user can combine in ways which suit his own cognitive style and problem solving methods. There is a danger that this could become tedious for the user after a period, and it is very ill-advised to create a design that requires people to do something they see as more onerous than their previous activity as they will quickly cease using the system.

This seems to offer scope for exploiting some of the techniques emerging from Artificial Intelligence research in the field of learning systems. If the scheduling system can be designed to observe the use made of the interactions in terms of frequency and sequences, then it may be able to identify higher level or more intelligent interactions combining several of the low level ones in a form consistent with the user's style. These could then be offered automatically to the user as additional facilities.

Another technical area worthy of further study is the generation of alternative schedules. The need for an operational researcher to work closely with the scheduler during the early stages of use to encourage him in this has already been referred to. This however is only a partial solution, and it is desirable to include in the system the capability to automatically generate tentative schedules for consideration by the user as good starting points from which he can develop a solution which is "optimal" in some sense. This has been achieved to a limited degree with some existing systems, but it is potentially a fertile area for further research.

Turning to the process of developing visual interactive simulation models, an issue of interest is the time and effort required to produce a working system. As indicated above the approach itself enables considerable reduction in both directions, but there remains scope for further worthwhile improvements.

To this end a number of software packages have been developed and are commercially available. These include SEE-WHY (Fiddy et al, [5]) and FORSSIGHT (Hollocks, [7]), both of which are based on the discrete event mechanism, and ICI's suite of DISIM (Displays, Interactions and SIMulation) programs which offers both the discrete event (Hill, [6]) and bar chart (Sykes, [10]) methods.

These are all in effect an enhanced FORTRAN, providing a library of subroutines covering many of the common features of visual interactive models, but allowing the user to easily add to these by coding additional ones himself. Such systems are a major advance for modelling efficiency but the building of systems remains manpower intensive and often lengthy.

Recent attention has been focussed on providing facilities which remove totally the need for any programming, at least for a limited class of problems, by extending the visual interactive approach to the model definition activity itself. With SEE-WHY this is being achieved through code generators, but with DISIM the preferred route has been to produce data driven interfaces using interactive editors to automate the design process.

Such developments may tempt users to develop models directly themselves without seeking the help of an operational researcher. The wisdom of such an approach is highly questionable, as modelling remains a skilled activity; the packages simplify the mechanics of constructing a model, but they do not provide the technical expertise to ensure the model is relevant, appropriate and requisite. The real advantage offered by the developments is that they enable the professional modeller to construct a model very rapidly indeed, often within hours and perhaps even jointly with the user at the terminal. This may be a model which the user can employ for real, or it may merely be a prototype/demonstration system used to help the communication between modeller and user. In either event the time savings in the overall modelling process can be great.

Although systems such as these are beginning to appear there remains considerable incentive to improve their flexibility, generality and performance.

Finally there is the possibility of life cycle software for batch chemical plants. Typically a batch plant is designed and built and then it proceeds through the iterative cycle of operation - modification - operation .... until it has exhausted its useful life. Discrete event simulation has long been used as an aid to tackling both the plant design and modification problems. Experience suggests that bar chart simulation is an effective aid to plant scheduling, but that the discrete event approach is particularly helpful at times of change to accelerate the learning experience of plant management.

This raises the interesting possibility of producing, for a given batch plant, a package of software relevant to all of the design, operation and modification activities. The package would be based on a single model for the plant with common data structures but perhaps different algorithms for the different activities, all wrapped up in suitable display and dialogue management facilities. Such a package would clearly be useful for each of the individual phases, but would also greatly facilitate the transition between phases.

9. CONCLUSION

Visual interactive modelling has rapidly become established as one of the prime growth areas of Operational Research with simulation applications well to the forefront. Experience gained from a number of scheduling projects clearly demonstrates that its use is not limited to existing application areas, which typically are one-off investigations. It can also be used beneficially on a much wider class of problems including routine day to day activities where decisions must be based not just on calculations but also on human skills. As Bell [1] suggests "visual interactive methods should be seen as problem-solving approaches, rather than model building techniques".

To be successful in exploiting simulation on this broader front will be very demanding in terms of:

- behavioural skills needed to manage the interpersonal and organisational issues likely to arise in the course of projects;
- technical skills needed to develop flexible, robust systems based on high quality interfaces and requisite models which may need to integrate several different techniques;
- creativity needed in the analysis and solution of novel problems, and the innovation of more powerful facilities for use in the models.

However, the potential rewards are high with the opportunity to impact significantly on decisions of significance to a company and thereby make a major contribution to improving business performance.

ACKNOWLEDGEMENT

The author wishes to thank the directors of ICI PLC for permission to publish this paper.

REFERENCES

[1] P.C. BELL. Visual interactive modelling in Operational Research. *J. Opl. Res. Soc*, vol 36, no 11, pp 975-982 (1985).

[2] H.C. BOWEN, R.F. FENTON, M.A.M. ROGERS, R.D. HURRION and R.J.F. SECKER. Interactive computing as an aid to decision makers. In: *Operational Research '78*, (Haley, K.B. Ed), North Holland. Amsterdam. pp 829-842 (1979).

[3] H.C. BOWEN, R.F. FENTON, D.E. CONNAUGHTON, M.J.F. FISHER and R.D. HURRION. An interactive aid to improve plant control. In: *IEE Conference Publication No. 172*. Third International Conference on Trends in On-line computer control systems, Sheffield, England, March 1979. pp 37-40 (1979).

[4] C.M. da SILVA, R.D. HURRION, W.H. SWANN and P.J. TOSNEY. A decision support system for the planning and control of complex interrelated manufacturing units. *IFORS Conference*, Hamburg (1981).

[5] E. FIDDY, J.G. BRIGHT and R.D. HURRION. See-why; interactive simulation on the screen. *Proc. Inst. Mech. Engrs.* C293/81. pp 167-172 (1981).

[6] S.R. HILL. The VUSIM visual discrete event simulation package. ICI Report No. IC06436 (1985).

[7] B.W. HOLLOCKS. Simulation and the micro. *J. Opl. Res. Soc.*, vol 35, no 4, pp 331-343 (1983).

[8] R.D. HURRION. An investigation of visual interactive simulation methods using the job-shop scheduling problem. *J. Opl. Res. Soc.*, vol 29, pp 1085-1093 (1978).

[9] D.W.T. RIPPIN. The present status and future prospects of computer-aided design and scheduling for batch processes. Technisch-Chemisches Labor, ETH Zentrum, Zurich Report SEG/R/81/107 (1981).

[10] J. SYKES. The BARNET barchart planning package. ICI Report No IC06437 (1985).